
Rule DAS100: VOLUME WITH WORST OVERALL PERFORMANCE

Finding: The identified volume had the worst overall performance during the entire measurement period. RULE DAS100 is quite similar to RULE DAS200; RULE DAS100 applies to **all** DASD devices, while RULE DAS200 applies only to DASD devices accessed by critical (or "loved one") workload.

Impact: The impact of this finding will depend upon the importance of the volume to overall system performance. If this is a critical volume, then this finding will have a HIGH impact. However, if the volume is accessed by low priority workloads, then this finding will have a LOW IMPACT or MEDIUM IMPACT.

Address spaces are retained in storage so long as they have uncompleted I/O operations. While the address spaces are in storage, they occupy page frames and may delay other address spaces. Additionally, the SRB time required to service I/O operations executes at a higher dispatching priority than a TCB, regardless of the dispatching priority of the TCB. Thus, there may be an overall system impact even though the volume may be accessed only by low priority workloads.

Logic flow: This is a basic rule finding; there are no predecessor rules.

Discussion: CPExpert determines the average device response time, by device type, for each measurement interval. A "device type" for this purpose is any unique device type (e.g., IBM-3380 or IBM-3390), with the device type modified to reflect whether the device is cached, has DLS or DLSE capabilities, or is a paging device.

The purpose of determining the average device response time, by device type, is the underlying principle that there is little point in analyzing a particular device if its response time is better than average. Rather, the most improvement potential resides with devices whose response time is worse than average.

CPExpert selects a device in each measurement interval for further analysis if the device response time exceeds the average for its device type.

CPExpert consolidates the channel path information (contained in the SMF Type 78 records) with the information contained in SMF Type 74 records. This information, along with the IOCP macro information, allows CPExpert

to build a model of the I/O configuration. This model includes all channel path, controller, and device utilization and queuing information. CPExpert then applies queuing formulae to the model to compute delays that occurred at significant parts of the model. The results from the model are associated with essential information describing the device response characteristics.

CPExpert constructs a frequency distribution of all devices whose response is worse than average, weighted by the number of I/O operations executed by the device. This yields a weighted measure of the potential performance improvement that might be achieved for each device. This frequency distribution is sorted descendingly, to yield an ordered list of the devices with the most improvement potential.

CPExpert selects the top devices from the ordered list of devices with the most improvement potential. The results from the model described above is extracted for each of these devices. Detailed information regarding the "worst" device is extracted and reported for each measurement interval. (RULE DAS190 reports summary results for the remaining top devices.)

Exhibit DAS100-1 provides a sample output resulting from the analysis. The VOLSER and device number of the "worst" performing device are identified in the narrative. Information is provided about the overall average I/O rate and the device utilization for the entire measurement period being analyzed.

As shown in Exhibit DAS100-1, CPExpert provides a summary for each measurement interval, showing the average I/O rate for the interval, and the **average delay time per second** during the interval (the total is shown as **I/O RESP** in Exhibit DAS100-1). The average delay time per second effectively reflects the percent of each second (shown in milliseconds) in which the associated delay occurred. The sum of the CONN, DISC, and PEND time is the average device busy time during an average second.

CPExpert analyzes each measurement interval and provides an indication of the major problem in the interval. The problems detected by CPExpert may be (1) excessive seeking, (2) missed RPS reconnect, (3) excessive connect, (4) missed cache hits, or (5) IOS queue. CPExpert determines the major cause of device response delays (simply dividing each potential delay by the total device response time). Additionally, the IOSQ time is divided by the total I/O service time.

The result from these calculations is evaluated to determine whether any area significantly predominates the I/O response time. If so, the respective area is listed as the major problem with the volume during the interval being analyzed. There are certain implications related to this analysis in each area of potential problem:

-
- **Excessive seeking.** CPEXpert computes an average seek delay time by subtracting the average latency and the missed RPS reconnect time from the total disconnect time. The remaining time is seek time for the volume. Excessive seeking is considered to be a major problem if the average seek time accounts for a significant portion of the device service time.
 - **Missed RPS reconnect.** CPEXpert creates a model of the I/O configuration, and assesses the probability that a missed RPS reconnect will occur (a M/M/C queuing model is used to estimate this probability). Missed RPS reconnect is determined to be a major problem if the number of missed RPS reconnects (multiplied by the device rotation time) accounts for a significant portion of the device service time.
 - **Excessive connect.** Excessive connect time is considered to be a major problem if the device connect time accounts for a significant portion of the device service time.
 - **IOS queue.** IOS queue time is considered to be a major problem if the IOS queue time accounts for a significant portion of the overall I/O response time. Note that IOS queue time is evaluated against the overall I/O response time. This is done because other problems (e.g., seeking) may cause excessive IOS queue time. The IOS queue time is unlikely to be resolved unless the underlying problem is addressed. Thus, IOS queue time must pass a more stringent test to be considered the major problem.
 - **Missed cache hits.** CPEXpert estimates the delay due to missed cache hits for cached devices. This estimation is done in the absence of RMF Cache Reporter data, and must be viewed with some caution. If the DISC time is less than the average latency, CPEXpert attributes the DISC time to missed cache hits.

If the DISC time is greater than the average latency, CPEXpert attributes the average latency to missed cache hits (the remainder of the DISC time is attributed to missed RPS reconnect and seek). If the estimated time attributed to missed cache hits accounts for a majority of the I/O delay time, CPEXpert concludes that the major problem with the volume is cache miss operations.

The rationale for this approach is that I/O operations to cached devices should not normally experience large average disconnect times. If they do experience large average disconnect times, either (1) many records are not found in the cache or (2) a large number of records require seek operations. If a cached volume experiences large average disconnect times, either the volume is a bad cache

candidate or else other volumes cause the cache to be ineffective for the volume.

The device may have no consistent problem. The problems may be concentrated in only a few measurement intervals and these few measurement intervals may dominate the performance characteristics of the volume. The worst of these intervals will be analyzed separately.

Additionally, it is possible that the device has no major problem for any interval. This condition most likely would occur if the guidance provided to CPEXpert is too restrictive. For example, if a large number of volumes are excluded from analysis (using the EXCLUDE option), the remaining volumes may have no particular problem. However, the logic is designed to select a "worst" volume, regardless of whether that volume actually has problems. CPEXpert tests for this condition and provides appropriate information.

Suggestion: There are no suggestions directly associated with this rule. Subsequent rules will analyze the device problems and attempt to determine the cause of poor performance.

RULE DAS100: VOLUME WITH WORST OVERALL PERFORMANCE

VOLSER RSA002 (device 0194) had the worst overall performance during the entire measurement period (13:59, 23JUL1991 to 16:59, 23JUL1991). This pack had an overall average of 17.2 I/O operations per second, was busy processing I/O for an average of 47% of the time, and had I/O operations queued for an average of 20% of the time. The following summarizes significant performance characteristics of VOLSER RSA002:

MEASUREMENT INTERVAL	I/O RATE	I/O RESP	CONN	DISC	PEND	IOSQ	MAJOR PROBLEM
13:59-14:14,23JUL1991	13.7	0.480	0.085	0.223	0.050	0.122	QUEUING
14:14-14:29,23JUL1991	16.5	0.581	0.102	0.270	0.049	0.160	QUEUING
14:29-14:44,23JUL1991	16.0	0.552	0.098	0.268	0.059	0.127	SEEKING
14:44-14:59,23JUL1991	16.5	0.572	0.101	0.289	0.054	0.128	SEEKING
14:59-15:14,23JUL1991	15.8	0.500	0.099	0.256	0.049	0.096	SEEKING
15:14-15:29,23JUL1991	16.9	0.665	0.096	0.288	0.073	0.208	QUEUING
15:29-15:44,23JUL1991	19.0	0.825	0.104	0.347	0.101	0.273	QUEUING
15:44-15:59,23JUL1991	18.5	0.737	0.103	0.350	0.053	0.231	QUEUING
15:59-16:14,23JUL1991	16.8	0.688	0.096	0.298	0.061	0.233	QUEUING
16:14-16:29,23JUL1991	19.9	0.793	0.108	0.354	0.058	0.273	QUEUING
16:29-16:44,23JUL1991	19.6	0.874	0.102	0.356	0.105	0.311	QUEUING
16:44-16:59,23JUL1991	17.5	0.669	0.098	0.320	0.061	0.190	QUEUING

VOLUME WITH WORST OVERALL PERFORMANCE

EXHIBIT DAS100-1